

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11) Publication number:

0 488 501 A2

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number: 91308228.5

(51) Int. Cl.<sup>5</sup>: G06F 9/46, G06F 3/06

(22) Date of filing: 09.09.91

A request for the exchange of the originally filed drawings against a new set of drawings has been filed pursuant to Rule 88 EPC. A decision on the request will be taken during the proceedings before the Examining Division (Guidelines for Examination in the EPO, A-V, 2.2).

(30) Priority: 26.10.90 US 605054

(43) Date of publication of application:  
03.06.92 Bulletin 92/23

(84) Designated Contracting States:  
DE FR GB IT

(71) Applicant: International Business Machines Corporation  
Old Orchard Road  
Armonk, N.Y. 10504(US)

(72) Inventor: Miro, Felix  
4928 S.W. 90th Avenue  
Cooper City, Florida 33328(US)

(74) Representative: Blakemore, Frederick Norman  
IBM United Kingdom Limited Intellectual  
Property Department Hursley Park  
Winchester Hampshire SO21 2JN(GB)

(54) Scheduling input/output operations in multitasking systems.

(57) In a multitasking data processing system, I/O requests to a disk drive are staged in holding queues from which they are transferred to a service queue. Requests in the latter queue are directly serviced on a FIFO basis by a device driver module running on the system. The system maintains a set of holding queues and an associated service queue separately for each physical drive in the system (hard file, floppy drive, etc.). Holding queues in each set are prioritised in accordance with base priorities of tasks, and I/O requests to disk drives are entered into associated holding queues having priorities corresponding to those of task threads for which such requests are originated. Prioritisation of the holding queues, and a starvation advancement process performed to advance "oldest" enqueued requests to higher priority holding queues, causes the requests to be presented to the disk drive in a sequence based in part on respective task priorities and in part on "fairness" servicing of "service starved" requests. A selection operation in respect to certain transfers from the holding queues to the service

queues orders selection of transferred requests on the basis of relative proximity of disk addresses in the queued requests to current positions of the read/write head assembly in the associated disk drive. Foregoing operations serve to improve processing throughput of all tasks. A service kernel of the operating system, which maintains the holding queues, places each newly issued disk I/O request in a selected holding queue associated with a designated disk drive. The selected queue is one having a priority associated with the base priority of the task thread for which the request was issued. A request issued for a foreground task may be placed in a higher priority holding queue than requests of like priority issued for other tasks. When entering a new request into a holding queue, the kernel performs a "starvation" check relative to selected holding queues. If a starvation condition is detected, the kernel performs the above-mentioned starvation advancement operation to transfer requests in the affected queue to a higher priority queue.

EP 0 488 501 A2

EP 0 488 501 A2

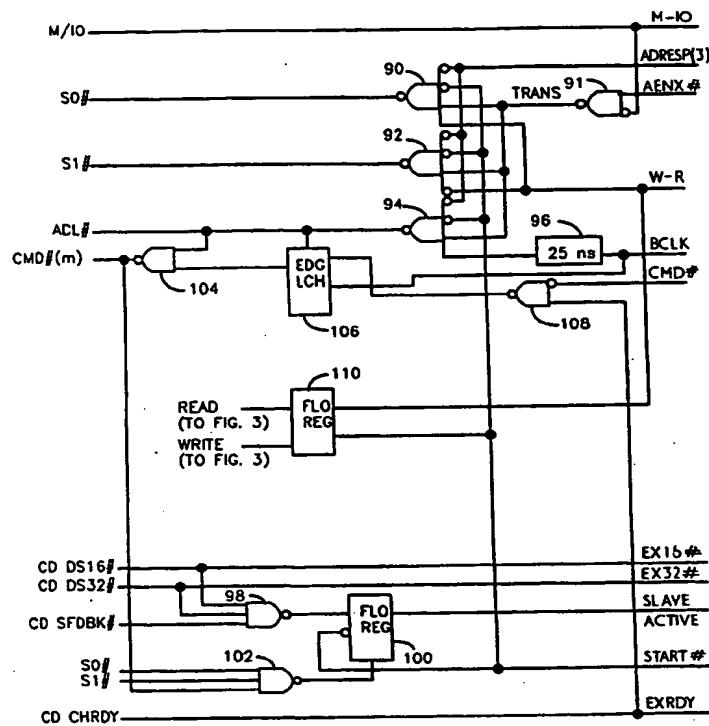


FIG. 5

Field Of The Invention

This invention relates to multitasking data processing systems; and particularly to the scheduling of input/output functions relative to disk storage devices in such systems.

Terminology

Terms which may be used frequently herein are defined below. Additional information regarding respective functions is contained in the publication IBM Operating System/2<sup>1</sup> Programming Tools and Information, Version 1.2 Programming Guide, 1989. Operating System/2 is also referred to herein as OS/2.

API (Application Programming Interface) - a programming interface between an operating system and application programs.

Device Driver - a program which performs operations to control a peripheral device; e.g. to direct transfers of data relative to a disk storage device.

Foreground Program - a program with which a user of a multitasking system is currently interacting.

Linked list queue - a linearly ordered queue of data items in which positions of consecutive items are indicated by pointer information in the items. In a single linked list item pointers point only to consecutive items in one direction, while in a double linked list queue item pointers point bidirectionally.

Multitasking - concurrent processing of applications or parts of applications in a data processing system.

RBA (Relative Block Address) - a number defining a storage location on a disk storage device. Magnitudes of such numbers are directly relatable to track and cylinder locations on the device.

Service Kernel - an OS/2 module, operating between API's and device drivers, to service requests initiated by program tasks.

Service Queue - a queue maintained by the operating system and a disk device driver, relative to an associated disk drive device, for ordering I/O functions performed by the device.

Task Request Packet - unit of request information passed from a program task to the operating system service kernel for initiating a required action such as I/O service.

Thread - a unit of execution within a task or process.

In contemporary multitasking data processing systems, operating systems, for instance, systems operating under operating system OS/2 versions 1.0, 1.1 and 1.2, input/output (I/O) operations relative to disk storage devices are scheduled in se-

quences which are essentially unrelated to the relative time urgencies of tasks waiting for such functions to be carried out.

For example, in OS/2 Versions 1.0, 1.1, and 1.2, I/O activities relative to a disk device are ordered in an "elevator" sequence in which I/O requests are arranged in an associated service queue in-a-positional sequence designed to allow servicing of all enqueued requests having targeted storage locations reachable while the read/write head of the disk is continuing in its present direction of motion. This is analogous to the familiar sequence followed by building elevators, wherein an elevator car moving upward continues upward to all selected floors, ignoring calls for service from floors which have been passed, and an elevator car moving downward car stops at all selected floors ignoring calls for service from higher floors which it has passed. Similarly, in prior art OS/2 systems disk I/O requests are enqueued alternately in sequences of increasing and decreasing associated RBA (relative block address) values representing their targeted disk storage locations.

In such systems, I/O requests directed to a specific disk drive are held in a single linked list service queue associated with that drive. The next request to be served is always the one located at the head end of the queue. While the R/W (read/write) head of the associated drive is moving in a direction of increasing RBA locations relative to the disk, requests designating RBA locations reachable in that direction (not yet passed by the head) are positioned in the queue in a sequence of increasing RBA's starting at the head end, and requests designating RBA locations reachable only after a reversal of R/W head direction are arranged further along the queue in a sequence of decreasing RBA's. When the operating system service kernel responsible for placing requests in the queue receives another request designating the same drive, it determines the present position and direction of movement of the R/W head (from the RBA's of the last 2 requests removed from the queue for service) and places the received request into the queue in a position for it to receive "elevator sequence" service in line with other enqueued requests.

Thus, if the R/W head is moving in the direction of increasing RBA disk locations, and a request at the head end of the service queue has an RBA representing a disk location reachable without reversing the direction of R/W head movement, all enqueued requests having such reachable RBA's are enqueued in a block having progressively ascending request RBA values starting at the head end of the queue. Any newly presented request

having an RBA representing a location reachable without a change in R/W head direction is inserted into the head end block so as to maintain the ascending order of request RBA's in the block. A new request having an RBA representing a location already passed by the head is placed in a second block in which the request RBA's are ordered in descending sequence. Thus, after all requests in the first block have been serviced, the direction of the R/W head will be reversed and requests in the second block will receive service consecutively.

Although such "elevator" ordering effectively guarantees that all requests will be serviced, it does not ensure efficient service since the priorities of tasks with which such requests are associated are not given any consideration. Thus, disk I/O requests having real-time association with a human heart operation stand to receive no better service than requests associated with continuation of a recreational game. Furthermore, since the rate of presentation of disk I/O requests in any system tends to increase in inverse relation to memory constraints (i.e. the smaller the system memory, the more frequent the access to disks) and in direct relation to the number of tasks concurrently processed, it is clear that in an efficiently utilised multitasking system with small memory capacity, the average delay in servicing disk I/O requests could become so large as to require system users to place undesirable constraints on their use of the system.

There is now provided according to the present invention a data processing system comprising a CPU, memory and disk storage subsystems, and means for scheduling I/O operations of said disk storage subsystem in accordance with I/O requests issued relative to tasks being processed by said CPU, characterised in that there are provided means to register a set of queues associated uniquely with a single disk drive in said storage subsystem, said queues serving to store I/O requests designating I/O operations to be performed by said associated disk drive; said set of queues comprising a service queue and at least one holding queue; device driver means coupled to said service queue for servicing requests stored in said service queue on a FIFO basis and responsive to each serviced request for dispatching an I/O operation designated by the request to be performed by said associated disk drive; means coupled to said holding queue or queues for receiving I/O requests newly issued relative to said processed tasks, and for storing said requests in said holding queue or queues in a sorted positional sequence in which said requests are positionally ordered in association with priority classes of tasks relative to which they are issued; and transfer means coupled to said holding queue or queues and said service

queue for selectively transferring requests from said holding queue or queues to said service queue in an order based at least in part on the task ordering of said requests in said holding queue or queues, whereby requests associated with tasks having a given priority class are transferred to said service queue before requests associated with tasks having a priority class lower than said given priority class.

There is further provided according to the present invention a data processing system comprising a CPU, memory and disk storage subsystems, and means for scheduling I/O operations of said disk storage subsystem in accordance with I/O requests issued relative to tasks being processed by said CPU, characterised in that there are provided means to register: a set of queues associated uniquely with a single disk drive in said storage subsystem, said queues serving to store I/O requests designating I/O operations to be performed by said associated disk drive; said set of queues comprising a single linked list service queue and plural double linked list holding queues having predetermined relative priorities for service, said holding queues including a plurality of queues having service priorities associated with priorities of tasks processed by said system and a FIFO holding queue having a service priority which is not associated with any task priority and is higher than the service priorities of the other holding queues; device driver means coupled to said service queue for servicing requests stored in said service queue on a FIFO basis and responsive to each serviced request for dispatching an I/O operation designated by the request to be performed by said associated disk drive; means coupled to said holding queues for receiving I/O requests newly issued relative to said processed tasks, each request containing a priority class indication of an associated task and an RBA (relative block address) value denoting a storage position on said disk drive relative to which a respective I/O operation is to be performed; means cooperative with said holding queue request receiving means for entering each said received request into a selected position in a holding queue other than said FIFO holding queue having a service priority corresponding to the task priority indicated in the respectively received request, said selected position being chosen as a function of RBA information in the respective received request and RBA information in requests presently held in the respective holding queue so as to maintain the requests in the respective queue positionally ordered with their RBA's forming a progressively increasing value sequence relative to one end of the respective holding queue; means coupled to said holding queues and said service queue, and responsive to an indication from the device driver

that the service queue is presently empty, for selectively transferring one or more requests from a highest priority non-empty one of said holding queues to said service queue; said last mentioned means transferring all requests from said FIFO holding queue to said service queue if the FIFO holding queue is not empty, and otherwise transferring a single request from a next highest priority non-empty holding queue if the FIFO holding queue is empty and another holding queue is not empty; said single request so transferred being selected from either the head or tail end of the respective next highest priority non-empty holding queue on the basis of a proximity determination made by said transferring means; said proximity determination being based on a comparison of the RBA values in the requests at the head and tail ends of the respective next highest priority non-empty holding queue with the RBA value of the request last served by said device driver; and starvation boosting means coupled to the means for entering said received requests for determining if a service starvation condition exists in either the holding queue into which said request is being entered or a lower priority holding queue, and for transferring all requests from the queue having said starvation condition to a selected higher priority holding queue; said selected higher priority queue being either the next highest priority empty one of said holding queues having priority less than said FIFO holding queue or being said FIFO holding queue if all holding queues of intermediate priority are not empty; said transferred requests when transferred to said FIFO holding queue being concatenated to the tail end of said queue if said FIFO holding queue is not empty when the transfer is made.

There is yet further provided according to the present invention a method of scheduling the handling of I/O requests relative to a disk drive in a multitasking data processing system comprising: entering said requests into a plurality of priority ordered holding queues associated with said disk drive; said holding queues having priorities for dequeuing service associated with classes of priority assignable to tasks processed by said system; said requests having indications of priority classes of related tasks and being entered into holding queues with associated priority; and transferring requests from said holding queues to a service queue associated with said disk drive, in the priority order of said holding queues; and dispatching requests from said service queue to said drive.

#### In the Drawings:

Figure 1 illustrates a typical prior art data processing system environment in which the present

invention may be advantageously applied.

Figures 2-4 are used to illustrate by specific example the prior art servicing of disk I/O requests in "elevator" sequence as broadly explained above.

Figure 5 schematically illustrates the present arrangement of holding and service queues for I/O requests designating a common disk drive.

Figure 6 illustrates a preferred arrangement of priority ranked holding queues in accordance with the present invention.

Figure 7 illustrates the placement of I/O requests in progressively ascending RBA sequence in the TC, RG and IDLE holding queues shown in Figure 6.

Figure 8 illustrates the ordering of I/O requests in the FIFO holding queue shown in Figure 6.

Figures 9 and 10 respectively illustrate request enqueueing and starvation advancement processes performed by the operating system service kernel relative to the holding queues shown in Figure 6.

Figure 11 illustrates the transfer process conducted by the service kernel, for transferring requests from the holding queues of Figure 6 to the service queue shown in Figure 5, in response to signals from the device driver which attends to requests in the service queue.

#### Detailed Description

##### 1. System Environment

Referring to Figure 1, a typical multitasking data processing system environment 1, in which the present invention can be used advantageously, comprises a microprocessor 2, a memory subsystem 3 (typically, comprising one or more banks of random access memory, and a not-shown direct memory access controller), a bus 4 and a disk storage subsystem 5 (typically comprising a controller not-shown and one or more disk drives). Associated with microprocessor 2 is Basic Input Output System (BIOS) "firmware" 2a of well known form, which may be stored in a read only memory (ROM) portion of the memory address space associated with subsystem 3. In a manner well-known in the art, bus 4 links elements 1-3, and 5 for exchange of information signals. Not-shown are user input devices typical of such systems (keyboard, mouse, etc.) and other system elements not considered relevant to an understanding of the present invention.

At system start-up, a multitasking operating system program 6, such as one of the existing versions of the OS/2 system, is loaded into memory subsystem 3 for managing utilisation of disk storage space in subsystem 5 and movement of information between that subsystem 5 and memory subsystem 3. The operating system operates to

provide dynamic transfer of programs and data between storage subsystem 5 and memory subsystem 3 in a manner permitting the microprocessor to be operated in a multitasking mode.

## 2. Handling of Disk I/O Requests By Prior Art Multitasking Operating Systems

Figures 2-4 illustrate the above-mentioned prior art "elevator" ordering of disk I/O requests in accordance with existing versions of the OS/2 system. In these OS/2 systems, application programs and utilities interface to the operating system through an application programming interface (API) 20. Disk I/O requests from the applications and utilities are passed through the API to a service kernel 22 of the operating system. Kernel 22 places each request into a service queue 24 associated with the disk drive device designated in the request. A separate service queue is maintained relative to each physical disk drive (hard file, floppy disk, etc.) in subsystem 5. Requests are fetched from queue 24 one at a time, in response to "device ready" indications from the device driver software module 26 associated with the respective drive, and dispatched to a not-shown controller of the designated disk drive (via BIOS).

Figures 3 and 4 provide a simplified illustration of how requests in prior art queue structure 24 are sortably positioned and serviced in "elevator sequence". The service queue 24 in these prior art systems is constructed as a single linked list queue into which requests are sortably placed by the service kernel, and from the head end of which requests are dispatched one at a time by the service kernel for action by the device. When service queue 24 is empty, the next request received by the service kernel relative to the associated device is placed at the head of the service queue. If the queue is not empty when a next request is received, the service kernel places that request at a position in the queue determined by the relation between the RBA designated by that request, the RBA's designated by requests already in the queue which imply the direction of movement of the associated disk drive's R/W head.

The received request is positioned in the service queue so as to maintain servicing of the enqueued requests in "elevator" sequence; i.e. so as to ensure that all enqueued requests having RBA's reachable by the R/W head without reversing its direction of movement will be serviced during movement of the head in its present direction and to ensure that enqueued requests having RBA's reachable in the reverse direction will be consecutively serviced after the direction of R/W head movement is reversed.

In the example shown in Figures 3 and 4, the

5 R/W head is assumed to be moving in the direction of increasing RBA's (since the RBA of the last dequeued request was 320 and the RBA of the previously removed request was 110), and the service queue is assumed to contain 5 unserviced requests. Two of the unserviced requests have RBA's with values of 430 and 1000, which represent storage locations reachable by the R/W head from its present position (at or beyond RBA 320) without a change of direction, and the other 3 requests have RBA's with values lower than 320 (specifically RBA values 126, 100 and 310) representing storage locations which cannot be reached without changing R/W head direction.

10 Notice that the requests having RBA's reachable in the present "ascending" direction of head movement are positioned in the queue with their RBA values forming a progressively ascending sequence "430, 1000", while the requests having RBA's not reachable without a change in R/W head direction are positioned in the queue with their RBA's arranged in a progressively descending sequence "310, 180, 126". Thus, if no more requests were placed in the queue the 2 requests closest to the head end would be dispatched to the disk while the R/W head is moving in its present ascending RBA direction, and the dispatch of the first of the three remaining requests would cause the R/W head to change direction and proceed progressively in the opposite direction; in which the head would reach the RBA storage locations of the last 3 requests consecutively.

15 20 25 30 35 40 Figures 3 and 4 show where in the service queue a new request with RBA value of 752 would be positioned, assuming that request is received when the queue is in the state shown in Figure 3. In Figure 3, the new request is shown above and outside the queue, and in Figure 4, the new request is shown in its appropriate "elevator ordered" position in the queue.

## 3. Present Disk I/O Scheduling - Overview

45 50 55 As shown in Figure 5, the present invention contemplates use of a modified service kernel 32 to maintain a set of (plural) holding queues 34 and a service queue 36 relative to each physical disk drive in the system. Service queue 36 differs from prior art service queue 24 in that its requests are serviced on a FIFO basis rather than in a sorted elevator sequence (in which requests are serviced in order unrelated to the sequence in which they are entered into the queue).

Upon receiving a device ready indication from driver 38, service kernel 32 dispatches a request from the head end of service queue 36, and if this empties the service queue, the kernel seeks to transfer a request from holding queues 34 to the

service queue. If at least one holding queue is then "ready" (not empty), the service kernel locates the highest priority such ready queue and transfers either one request or all requests from that queue to the service queue (in a manner described in further detail below). If the holding queues are all empty at the time, the service kernel will advance the next received request directly to the service queue.

#### 4. Present Disk I/O Scheduling - Specifics

Figures 6-11 show the organisation and usage of a preferred configuration of holding queues in accordance with the invention. The holding queues and their task priority associations are shown in Figure 6. Ordering of requests within individual holding queues is shown in Figures 7 and 8. Operations performed by the service kernel to enter newly received requests into holding queues are shown in Figure 9. Service starvation and priority boosting operations performed by the service kernel relative to the holding queues are shown in Figure 10. Figure 11 shows operations performed by the service kernel for conditionally transferring requests from the holding queues to the service queue.

Referring to Figure 6, the holding queues are double linked lists (described further in Figures 7 and 8 having the following designations, in descending priority order: FIFO, TC\_4, TC\_3, TC\_2, TC\_1, RG\_4, RG\_3, RG\_2, RG\_1, and IDLE. The FIFO holding queue receives requests only from lower priority queues during a starvation priority boosting procedure described later with reference to Figure 10. The other holding queues have priorities corresponding to task priority classes and receive new requests associated with tasks having corresponding priorities. Queues having priorities intermediate those of the FIFO and IDLE queues are also eligible to receive requests from lower priority queues during the starvation boosting process to be described.

In OS/2 systems, tasks have priority classifications assigned by their programmers. Up to 96 distinct classes of task priority are assignable, falling into three major groups as follows: a group of 32 time critical task classes (TC31 - TC00, in descending priority order), a group of 32 regular/intermediate priority task classes (RG31 - RG00, in descending priority order), and a group of 32 lowest/idle priority task classes (IDLE31 - IDLE 00, in descending order).

Queues TC\_n (n=1-4) receive new requests associated with tasks classified TC31-TC00 in the four subgroups shown in Fig. 6. Thus, TC\_4 receives new requests associated with task priorities TC24 through TC31; TC\_3 receives new re-

quests associated with task priorities TC16-TC23; TC\_2 receives requests associated with priorities TC08-TC15; and TC\_1 receives requests associated with TC00-TC07. Requests associated with foreground tasks with time critical classification are given favoured treatment in relation to requests associated with non-foreground tasks of the same priorities. Thus, new requests associated with foreground tasks in any of the classes TC00-TC31 (TC/fgnd in Fig. 6) are all directed to TC\_4 (the highest priority time critical queue).

Similarly, queues RG\_n (n=1-4) receive new requests associated with tasks in the intermediate (regular priority) range RG31-RG00, in four major sub-groupings shown in Fig. 6. RG\_4 receives new requests associated with task priorities RG24 through RG31; RG\_3 receives requests associated with priorities RG16-RG23; RG\_2 receives requests associated with RG08-RG15; and RG\_1 receives requests associated with RG00-RG07. Requests associated with foreground tasks with intermediate/regular classification are given favoured treatment. Thus, new requests associated with foreground tasks in any of the classes RG00-RG31 (RG/fgnd in Fig. 6) are all directed to RG\_4 (the highest priority RG queue).

Figure 7 shows how requests are positioned in each of the holding queues other than the FIFO queue. As shown in Figure 7 each queue is constructed as a double linked list referred to an anchor table 60. Anchor 60 is located at an address in system memory which remains fixed at least during a given processing session. When the queue contains at least 2 requests, anchor 60 contains pointers for locating requests the head and tail end requests 61 and 62. Requests at intermediate positions are suggested at 63, and the double linkage between successive request positions is suggested at 64. It is understood that request positions 61-63 may be arbitrarily located in system memory. Requests are positioned in the queue with their RBA values progressively increasing or ascending, in the direction of the tail end, as shown in the drawing.

In addition to providing indications for quickly (directly) locating head and tail end requests, the anchor provides other key information for quickly ascertaining the queue status (without having to locate and examine the individual request packets). This includes a MaxWaitTime function representing a timeout factor for starvation service (discussed later re Fig. 10), a TimeStamp function which when the queue receives a request after being empty is the Present Time Of Day increased by the MaxWaitTime factor, and a BoostQ function which may be used to indicate a queue to which to shift contents of the present queue when a starvation condition is detected as described later. In addition

to the individual queue anchors, an I/O structure table (IOStruc) is maintained in memory relative to each physical disk drive for indicating the set of all queues associated with that drive, processing status of the associated device driver, the RBA position of the R/W head of the drive after it completes handling of the request most recently dispatched from the service queue, and the identity of the highest priority non-empty holding queue.

As shown in Figure 8, the FIFO holding queue is also constructed as a double linked list queue, with an anchor indicating key parameters including locations of end requests. The FIFO queue receives requests only from lower priority queues as a result of starvation handling discussed later (Figure 10). Since requests are transferred to this queue in blocks during starvation boosting, requests in the FIFO queue when there are any consist of one or more blocks within each of which requests (when there are more than one) have increasing RBA's (relative to the head end).

Operations performed by the operating system service kernel, upon receiving an I/O request directed to a given disk drive M, are indicated in Figure 9. In the example shown, the request is assumed to be associated with a task N having priority class TC07. From information in the request packet, the service kernel ascertains the targeted drive and the associated task priority (step 70). Next, the kernel ascertains if the associated task is or is not being processed in the operating system foreground (step 72), again from information in the request packet. If the task is in the foreground, operation 74 is performed to place the request in holding queue TC\_4M (the highest priority time critical holding queue associated with drive M). If the task is not in the foreground, operation 76 is performed to place the request in queue TC\_1M (the time critical holding queue associated ordinarily with tasks in classes TC00 through TC07).

As noted earlier, requests in the priority classed holding queues are ordered positionally in each queue with their RBA's increasing towards the tail end of the queue. Accordingly, it is understood that such ordering is maintained in request entry operations 74 and 76. Upon completing operation 74 or 76, the service kernel performs the starvation check and selective boosting process described next with reference to Figure 10.

In the starvation check process (operations 80-90 in Figure 10), the service kernel first determines at 80 if the holding queue in which it has just entered a request (assumed to be TC\_4M relative to the example of Figure 9) was empty before the entry (by examining status information in the respective queue anchor). If the queue was previously empty, no further starvation checking is needed relative to this queue, and a new

TimeStamp (equal to the Present Time plus the MaxWaitTime factor for this queue) is calculated and placed in the respective queue anchor (operation 82). At such times, the kernel proceeds with operations beginning at 84 to try to find a next lower priority queue needing starvation service.

If the kernel finds in operation 80 that the queue in which a request was just entered was not previously empty, it compares the respective anchor TimeStamp with the Present Time to determine if the timeout factor associated with the MaxWaitFunction has elapsed (operation 86). If the TimeStamp is greater than the present time, the timeout has not ended, so no further checking is needed, and the kernel proceeds to operation 84 for checking lower priority queues. If the TimeStamp is less than Present Time (Y at step 86), this means that the MaxWaitTime timeout has been passed and operation 87 is performed to check the individual request packets in the queue for starvation conditions.

This further checking of individual request packets in the queue is needed because the TimeStamp initially recorded is associated with entry of a request into a previously empty queue (operation 82) and the request so entered could have been transferred at any time to the service queue (refer to discussion of Figure 11 below) without the associated TimeStamp being changed. Thus, the anchor TimeStamp might not reflect the true status of the queue (i.e. the age of its "oldest" request). Each request packet has an individual TimeStamp entry which is the sum of the Time of its entry and the MaxWaitTime. If these individual TimeStamps are all not less than Present Time (N exit at operation 87) the queue is not starved and the kernel proceeds to operation 84 (at such times, the kernel sets the anchor TimeStamp value equal to that of the oldest/smallest individual TimeStamp value). If any of the individual TimeStamps is less than Present Time (Y exit at operation 87), its timeout has passed and the queue is considered to be in a "starved" condition. The service kernel then performs the boosting process of operations 92 and higher to transfer the contents of the starved queue to a next higher priority "eligible" holding queue.

An "eligible" holding queue other than the FIFO queue (which is "always eligible") is one which is either empty or presently contains requests whose RBA's are all either higher than the RBA of the tail end request in the starved queue or lower than the RBA of the head end request in that queue. Furthermore, the operation 92 search for a "next higher priority eligible queue" begins at a "boost target queue" designated in the anchor of the starved queue (this permits the system to skip over adjacent priority queues when boosting a

starved queue).

Depending upon whether the next higher priority eligible queue found in operation 92 is the FIFO holding queue or another (intermediate priority) queue (Y or N, respectively at decision 94), the system transfers (boosts) the contents of the starved queue respectively to the FIFO queue or other found queue (operation 96 or 98, respectively) and resets the starved queue to empty. If the eligible queue to which the requests in the starved queue are transferred was not empty before the transfer, the starved queue requests are concatenated in a block to one end of that eligible queue. If the respective eligible queue is the FIFO queue, the starved queue requests are concatenated to the tail end of the FIFO queue. If the respective eligible queue is other than the FIFO queue, the requests in the starved queue are concatenated to whichever end of the eligible queue is appropriate for keeping the sequence of RBA's in the respective eligible queue progressively increasing from head end to tail end after the transfer (per Fig. 7).

It should be understood that since the queues are double linked lists, block transfer 96 or 98 can be accomplished simply by transferring head and/or tail pointer information from the anchor of the starved queue to the anchor of the selected eligible queue, and if necessary, modifying request packets at a position of concatenation in the eligible queue to cross-link to each other. If the eligible queue is empty before the transfer, head and tail pointers are transferred from the starved queue's anchor to the eligible queue's anchor and no request packets are modified.

If the located eligible queue is not empty before the transfer, and it is the FIFO queue, the tail pointer in the starved queue anchor replaces the tail pointer in the FIFO queue's anchor, and request packets at the position of concatenation (what was the tail end of the FIFO queue before the transfer) are modified to cross-link to each other (i.e. the request packet which was at the tail end of the FIFO queue before the transfer and the packet which was at the head end of the starved queue are so modified).

If the located eligible queue is other than the FIFO queue, and was not empty before the transfer, a selected end pointer (head or tail) in the starved queue's anchor replaces a respective end pointer in the eligible other queue's anchor, and the request packets which become newly positioned next to each other, at the position of concatenation in the eligible other queue, are modified to cross-link to each other. Thus, if starved requests are concatenated to the head end of the eligible other queue (all starved request RBA's less than RBA's of all requests previously in the other queue), the

5 head end pointer in the starved queue's anchor becomes the new head end pointer in the anchor of the eligible other queue, and the requests which previously were the tail end and head requests in the starved and other queue are modified to cross-link. Conversely, if starved requests are concatenated to the tail end of the other queue (all starved request RBA's greater than those of all requests previously in the other queue), the tail end pointer in the starved queue's anchor replaces the previous tail end pointer in the other queue's anchor and the respective previous head end request of the starved queue and tail end request of the other queue are modified to cross-link.

10 The reason for maintaining progressively increasing RBA's, when concatenating starved requests to a non-empty eligible queue of intermediate priority, is to allow for selection of a request from one end of the respective queue to the service queue in accordance with a head proximity determination (refer to following discussion of Figure 11).

15 Figure 11 shows how requests are transferred from holding queues to the service queue. Upon 25 receiving a "drive ready" indication from the device driver relative to a disk drive (at 110), which is given by the latter when the drive is ready for an I/O operation, the service kernel dispatches (dequeues) the head end request from the associated service queue for initiating a respective I/O operation. If the service queue is empty after dispatching that request (Y at decision 112), the kernel proceeds to try to find the highest priority ready (non-empty) holding queue and if it finds such to transfer one or more requests from it to the service queue (operations 120 and higher).

20 30 35 40 45 If not all holding queues are empty, the service kernel determines if the FIFO queue is not empty; i.e. if it is the highest priority ready queue. If the FIFO queue is ready (Y at decision 122), all of its requests are transferred in a block to the service queue in their existing RBA order and the FIFO queue is reset to empty (operation 124). If the highest priority ready queue is other than the FIFO queue (N at decision 122), a single request is transferred from the head or tail end of that other queue to the service queue and indicators in the other queue are modified to (indicate a new head or tail end request position (operation 126).

50 55 Selection of the head or tail end request in operation 126 is made according to a proximity determination relative to the present position of the read/write head in the associated disk drive. The "end RBA's" in the head and tail end requests ("end RBA's" are the sums of RBA's and designated transfer block lengths in respective request packets) are compared to the end RBA of the request last dispatched to the disk drive (in opera-

tion 110), the value of the latter end RBA being recorded in the associated IOStruc table and effectively serving as an indication of the present position of the read/write head. The head or tail end request with the RBA closest to the read/write head position is transferred to the (head) of the service queue.

In summary the advantages of the invention are achieved by providing a set of priority ordered request holding queues relative to each physical disk drive in a processing system (hard drive, floppy drive, etc.). Requests directed to a given drive are entered into holding queues having service priorities corresponding to priority classes of tasks with which such requests are associated. Each queue is a double linked list queue, and requests placed in each are positioned to order their associated RBA (relative block address) values (which define storage locations on the drive respectively targeted for I/O action) in a progressively increasing sequence relative to one end of the queue.

When a last request on the service queue is dispatched to the disk drive, a request is transferred from a highest priority holding queue "ready" to be served (not empty) to the service queue. When the highest priority "ready" queue is other than the highest priority holding queue (herein termed the FIFO holding queue), a single request is transferred from a selected from one of the two ends of the respective ready queue. The selection is based on a "proximity" determination between the RBA's of the 2 end requests and the present position of the disk read/write head (the latter inferred from the RBA of the request last dispatched from the service queue).

In a preferred embodiment of the invention, a set of 10 holding queues is assigned to each physical disk drive. In each set, 9 of the 10 queues are ordered in association with 9 different groups of priority classes assignable to tasks. The remaining queue is a highest priority FIFO holding queue mentioned above. The FIFO holding queue receives requests only from lower priority queues during a specialised "starvation boosting" process considered unique to the present invention. In that process, when a determination is made that an "oldest" request in a holding queue has been enqueued for longer than a predetermined "starvation time", all requests in that queue are transferred in a block to a next higher priority holding queue which is either the next higher priority empty holding queue or the FIFO holding queue if all intermediate priority holding queues are not empty. If the holding queue is not empty prior to a request transfer, the transferred requests are concatenated to the request at the tail end of the FIFO queue.

The foregoing queues are managed by a ser-

vice kernel of the operating system. When the disk drive is ready for an I/O operation, the service kernel is alerted by the associated device driver, and dispatches the next request on the service queue (the requests are dispatched in FIFO sequence from the that queue). If the service queue is empty after a request is dispatched, and the holding queues are not empty, the service kernel transfers one or more requests from the highest priority non-empty holding queue to the service queue. If the FIFO holding queue is not empty, all the requests in that queue are transferred in a block to the service queue. If the FIFO queue is empty when the service queue becomes empty, and any lower priority holding queue is not then empty, a single request is transferred to the service queue from the non-empty holding queue next highest in priority relative to the FIFO holding queue. If that next highest priority queue contains more than one request, RBA's designated in the requests at the head and tail end of that queue are compared to the current R/W head position and the request having an RBA closest to that position is selected. In the preferred embodiment, the compared request RBA's are "end RBA's" (the RBA specified in the request increased by the file length specified in the request) and the request selected is the one having the closest end RBA value.

Requests entered into each holding queue other than the FIFO holding queue are positionally ordered to maintain a progressively ascending sequence of request RBA's beginning at one end of the queue (the head end in the embodiment to be described). Thus, requests in the FIFO holding queue consist of one or more blocks of requests transferred from lower priority holding queues, requests in each block positioned to maintain progressive linear ordering of their RBA's (e.g. ascending order in the embodiment to be described).

Each disk I/O request presented to the service kernel of the operating system contains a priority indication corresponding to the base priority class of the associated task. The service kernel places each request into a selected one of the holding queues in the set of holding queues associated with the targeted disk drive. The selected queue is one having a priority either corresponding to or higher than that of the priority indication in the task (but not the FIFO holding queue). I/O requests associated with foreground tasks are placed in a holding queue of corresponding or higher priority, whereas I/O requests associated with non-foreground tasks are placed only in a holding queue of corresponding priority. Thus, relative to tasks of a given priority class, I/O requests associated with foreground tasks tend to receive more favourable service than requests associated with non-foreground tasks.

A feature of the present hierarchical arrangement of priority holding queues, with linear RBA positional ordering in each queue, is that requests are effectively prioritised for service in accordance with priority classes of associated tasks and selectable for service based on relative proximity of end RBA's to the disk R/W head; so that only two the two requests most quickly locatable in any queue need be examined for selection.

All of the foregoing features in combination serve to improve task processing throughput significantly, relative to comparably configured systems employing the prior art "elevator sequence".

In the preferred embodiment, the holding queues in each set with service priority ranking intermediate the highest (FIFO) and lowest (IDLE) levels, include a subset of time critical (TC) queues and a subset of regular (RG) priority queues. The TC queues are ranked higher in priority for device driver service than the RG queues. The TC and RG queues respectively receive requests associated with tasks having time critical and regular non-critical priority classifications. Requests sent to the IDLE holding queue are those associated with tasks having non-critical priority classifications lower than regular non-critical.

The foregoing arrangement of prioritised holding queues, in combination with the "starvation" handling process mentioned above and the request RBA ordering and head/tail proximity selection as characterised above, appear to provide optimum efficiency of request servicing with associated optimum improvement in overall task processing throughput. With this arrangement, throughput increases on the order of 3 to 5 percent are obtainable (relative to comparably configured systems using the "elevator" arrangement of I/O request servicing described above).

#### Claims

1. A data processing system comprising a CPU, memory and disk storage subsystems, and means for scheduling I/O operations of said disk storage subsystem in accordance with I/O requests issued relative to tasks being processed by said CPU, characterised in that there are provided means to register

a set of queues associated uniquely with a single disk drive in said storage subsystem, said queues serving to store I/O requests designating I/O operations to be performed by said associated disk drive; said set of queues comprising a service queue and at least one holding queue;

device driver means coupled to said ser-

5

vice queue for servicing requests stored in said service queue on a FIFO basis and responsive to each serviced request for dispatching an I/O operation designated by the request to be performed by said associated disk drive;

10

means coupled to said holding queue or queues for receiving I/O requests newly issued relative to said processed tasks, and for storing said requests in said holding queue or queues in a sorted positional sequence in which said requests are positionally ordered in association with priority classes of tasks relative to which they are issued; and

20

transfer means coupled to said holding queue or queues and said service queue for selectively transferring requests from said holding queue or queues to said service queue in an order based at least in part on the task ordering of said requests in said holding queue or queues, whereby requests associated with tasks having a given priority class are transferred to said service queue before requests associated with tasks having a priority class lower than said given priority class.

25

2. A system in accordance with claim 1, wherein:

30

said device driver means includes means for detecting when said service queue is empty and means for providing an empty indication associated with that condition; and

35

said transfer means includes means for detecting when a holding queue is not empty, and means responsive to said service queue empty indication when said holding queue is not empty for transferring at least one request from said holding queue to said service queue in said task priority based order.

40

3. A system in accordance with claim 2, wherein said processed tasks have at least a first priority class and at least a second priority class representing a lower priority class than said first priority class; and there are holding queues including:

45

a highest priority holding queue;

50

at least one first priority holding queue; and

55

at least one second priority holding queue; and wherein:

said means for storing said received requests is adapted to store requests associated with tasks having said first and second priority classes respectively in said first priority holding queue or queues and said second priority holding queue or queues; and wherein:

said transfer means is responsive to each said service queue empty indication to transfer a request from said highest priority holding queue to said service queue if the highest priority holding queue is not empty, or to transfer a request from a first priority holding queue to the service queue if the highest priority queue is empty and the first priority queue is not empty, or to transfer a request from a second priority holding queue to the service queue if the highest priority and first priority queues are empty and the second priority queue is not empty.

4. A system in accordance with claim 3, further including starvation processing means for detecting service starvation conditions in said second and first holding queues, said conditions occurring when a request has been held in the respective holding queue for more than a predetermined time and for transferring requests from the queue in which said condition is detected to a selected higher priority queue.

5. A system in accordance with claim 2, wherein the priority classes of said processed tasks comprise at least first and second groups of time critical priority classes and at least first and second groups of regular priority class, said classes having the following relative priority ranking in descending order of priority:

first time critical;

second time critical;

first regular;

second regular;

and wherein the holding queues comprise a highest priority holding queue, first and second time critical priority holding queues and first and second regular priority holding queues having the following relative priority rankings for service in descending order:

highest priority,

first time critical priority,

5

10

15

20

25

30

35

40

45

50

55

second time critical priority,

first regular priority,

second regular priority;

and wherein said means for storing said received requests in said at least one holding queues is adapted to store said requests in accordance with the following schedule:

store requests associated with foreground tasks having priority classifications in either said first time critical or second time critical class groups in said first time critical holding queue;

store requests associated with foreground tasks having priority classifications in either said first regular or second regular class groups in said first regular holding queue; and

store requests associated with non-foreground tasks having priority classifications in said first time critical, second time critical, first regular and second regular class groups respectively in said first time critical, second time critical, first regular and second regular holding queues;

and wherein said transfer means is responsive to each said service queue empty indication to transfer a request from the highest priority non-empty holding queue to said service queue when at least one of said holding queues is not empty.

6. A system in accordance with claim 5, wherein each of said time critical and regular holding queues is constructed as a double linked list queue.

7. A system in accordance with claim 6, wherein:

each of said requests contains RBA information denoting a relative block storage address in said associated disk drive relative to which an I/O operation is to be performed;

said means for storing said received requests in said holding queues is operative when entering requests into said time critical and regular holding queues to position the entered request relative to other requests in the queue so as to maintain the RBA values of all the requests in the queue positionally or-

dered in a progressively ascending numerical sequence relative to one end of the respective queue;

5 said device driver indicating means provides an RBA indication with said service queue empty indication, said RBA indication indicating the RBA value of the request last removed from said service queue and also indirectly indicating the current position of the read/write (R/W) head assembly in said associated disk drive; and

10 15 said means for transferring requests from said holding queues to said service queue is operative when removing a request for that purpose from any of said time critical or regular holding queues to select the request so removed from one or the other end of the respective queue, the selected request being chosen on the basis of a proximity calculation determination such that the request selected is the one having an RBA value closest to the current position of said R/W head assembly in said associated disk drive as indicated by said indication of said RBA value of the request last removed from said service queue.

20 25 30 35 40 45 50 55 8. A data processing system comprising a CPU, memory and disk storage subsystems, and means for scheduling I/O operations of said disk storage subsystem in accordance with I/O requests issued relative to tasks being processed by said CPU, characterised in that there are provided means to register:

a set of queues associated uniquely with a single disk drive in said storage subsystem, said queues serving to store I/O requests designating I/O operations to be performed by said associated disk drive; said set of queues comprising a single linked list service queue and plural double linked list holding queues having predetermined relative priorities for service, said holding queues including a plurality of queues having service priorities associated with priorities of tasks processed by said system and a FIFO holding queue having a service priority which is not associated with any task priority and is higher than the service priorities of the other holding queues;

device driver means coupled to said service queue for servicing requests stored in said service queue on a FIFO basis and responsive to each serviced request for dispatching an I/O operation designated by the request to be performed by said associated

disk drive;

means coupled to said holding queues for receiving I/O requests newly issued relative to said processed tasks, each request containing a priority class indication of an associated task and an RBA (relative block address) value denoting a storage position on said disk drive relative to which a respective I/O operation is to be performed;

means cooperative with said holding queue request receiving means for entering each said received request into a selected position in a holding queue other than said FIFO holding queue having a service priority corresponding to the task priority indicated in the respectively received request, said selected position being chosen as a function of RBA information in the respective received request and RBA information in requests presently held in the respective holding queue so as to maintain the requests in the respective queue positionally ordered-with-their RBA's forming a progressively increasing value sequence relative to one end of the respective holding queue;

means coupled to said holding queues and said service queue, and responsive to an indication from the device driver that the service queue is presently empty, for selectively transferring one or more requests from a highest priority non-empty one of said holding queues to said service queue; said last mentioned means transferring all requests from said FIFO holding queue to said service queue if the FIFO holding queue is not empty, and otherwise transferring a single request from a next highest priority non-empty holding queue if the FIFO holding queue is empty and another holding queue is not empty; said single request so transferred being selected from either the head or tail end of the respective next highest priority non-empty holding queue on the basis of a proximity determination made by said transferring means; said proximity determination being based on a comparison of the RBA values in the requests at the head and tail ends of the respective next highest priority non-empty holding queue with the RBA value of the request last served by said device driver; and

starvation boosting means coupled to the means for entering said received requests for determining if a service starvation condition exists in either the holding queue into which

said request is being entered or a lower priority holding queue, and for transferring all requests from the queue having said starvation condition to a selected higher priority holding queue; said selected higher priority queue being either the next highest priority empty one of said holding queues having priority less than said FIFO holding queue or being said FIFO holding queue if all holding queues of intermediate priority are not empty; said transferred requests when transferred to said FIFO holding queue being concatenated to the tail end of said queue if said FIFO holding queue is not empty when the transfer is made.

9. A method of scheduling the handling of I/O requests relative to a disk drive in a multitasking data processing system comprising:

entering said requests into a plurality of priority ordered holding queues associated with said disk drive; said holding queues having priorities for dequeuing service associated with classes of priority assignable to tasks processed by said system; said requests having indications of priority classes of related tasks and being entered into holding queues with associated priority; and

transferring requests from said holding queues to a service queue associated with said disk drive, in the priority order of said holding queues; and

dispatching requests from said service queue to said drive.

10. A method in accordance with claim 9 including:

intermittently performing a starvation check process relative to one or more of said holding queues; and

upon detecting a service starved condition in a checked holding queue, transferring all requests from said checked queue to a higher priority holding queue.

11. A method in accordance with claim 9 including:

favouring the handling of requests associated with foreground tasks over requests associated with non-foreground tasks of equal priority class by entering the requests associated with foreground tasks into higher priority said holding queues.

12. A method in accordance with claim 10 including:

organising said holding queues as double linked lists;

arranging requests entered into said holding queues in a sorted sequence wherein RBA (relative block address) contained in said requests and designating storage locations on said disk drive are progressively ordered in an ascending sequence in each said holding queue; and

when transferring a request from one of said holding queues to said service queue, selecting a request from either the head end or tail end of the respective holding queue on the basis of the relative proximity of the RBA location designated by the selected request to the present RBA location of the read/write head of the disk drive.

13. A method in accordance with claim 12 wherein:

the present location of said read/write head is inferred from the RBA of a request last dispatched from said service queue.

14. The method of claim 13 wherein said RBA of the request last dispatched from said service queue is an end RBA determined by summing an initial RBA function in the dequeued request and a block length parameter specified in the same request.

15. The method of claim 14 wherein said end RBA is stored in a key parameter table maintained relative to said service queue and the associated holding queues.

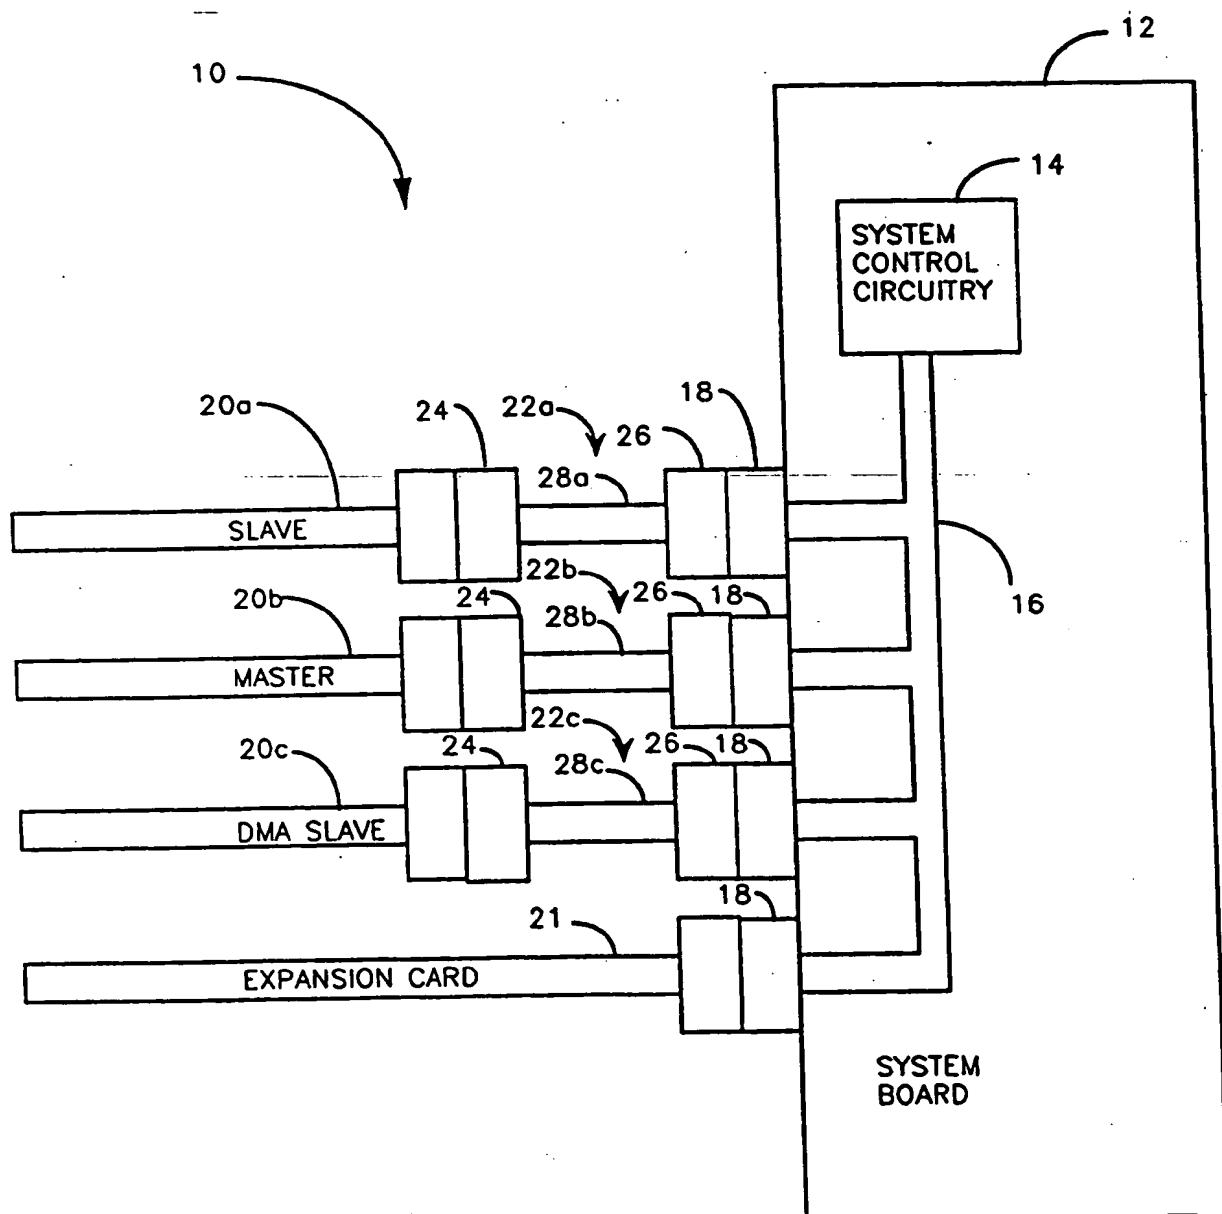


FIG. 1

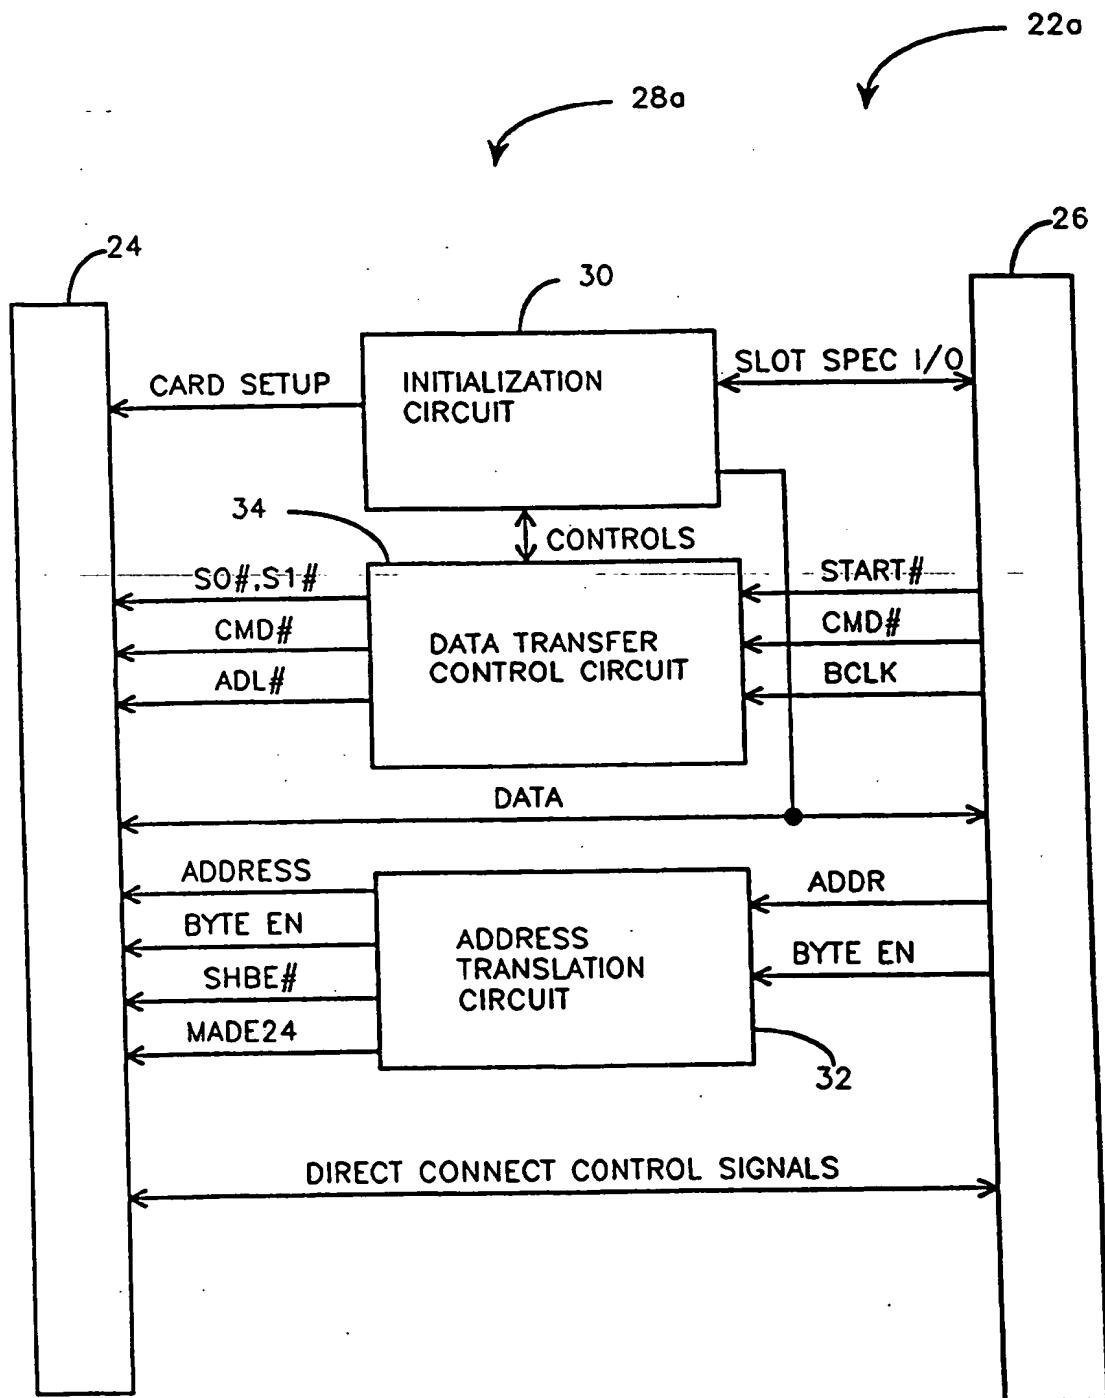


FIG. 2

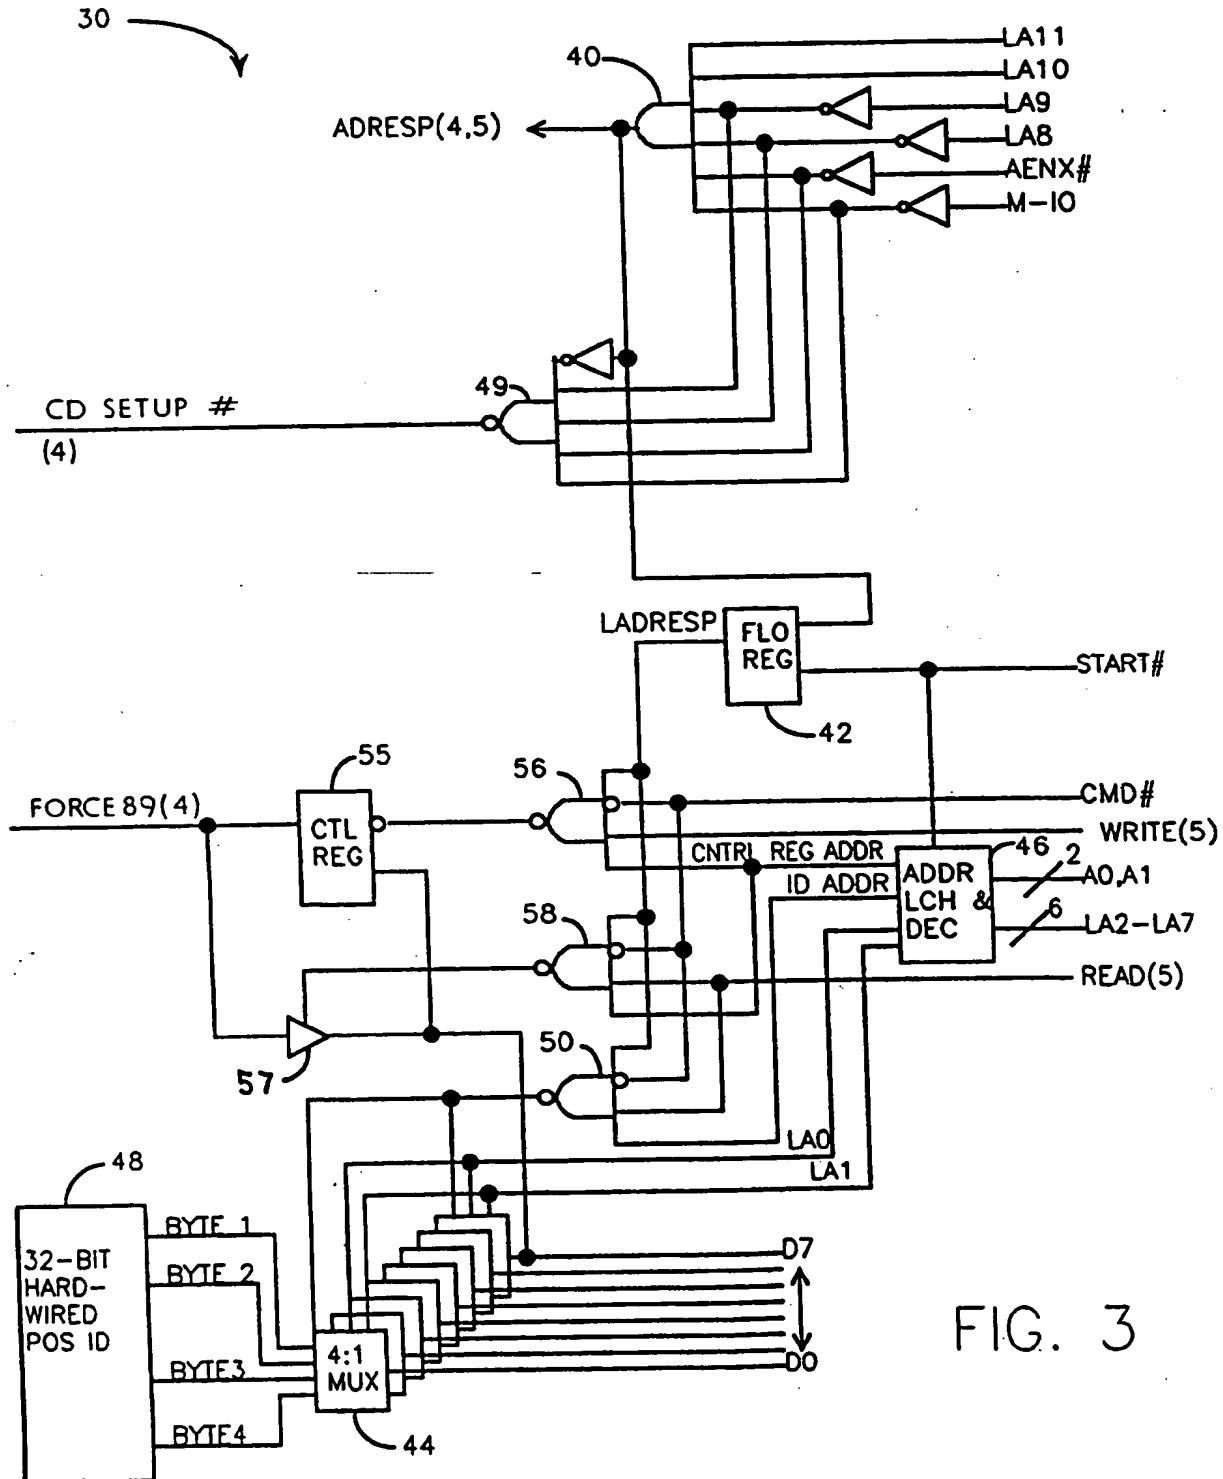
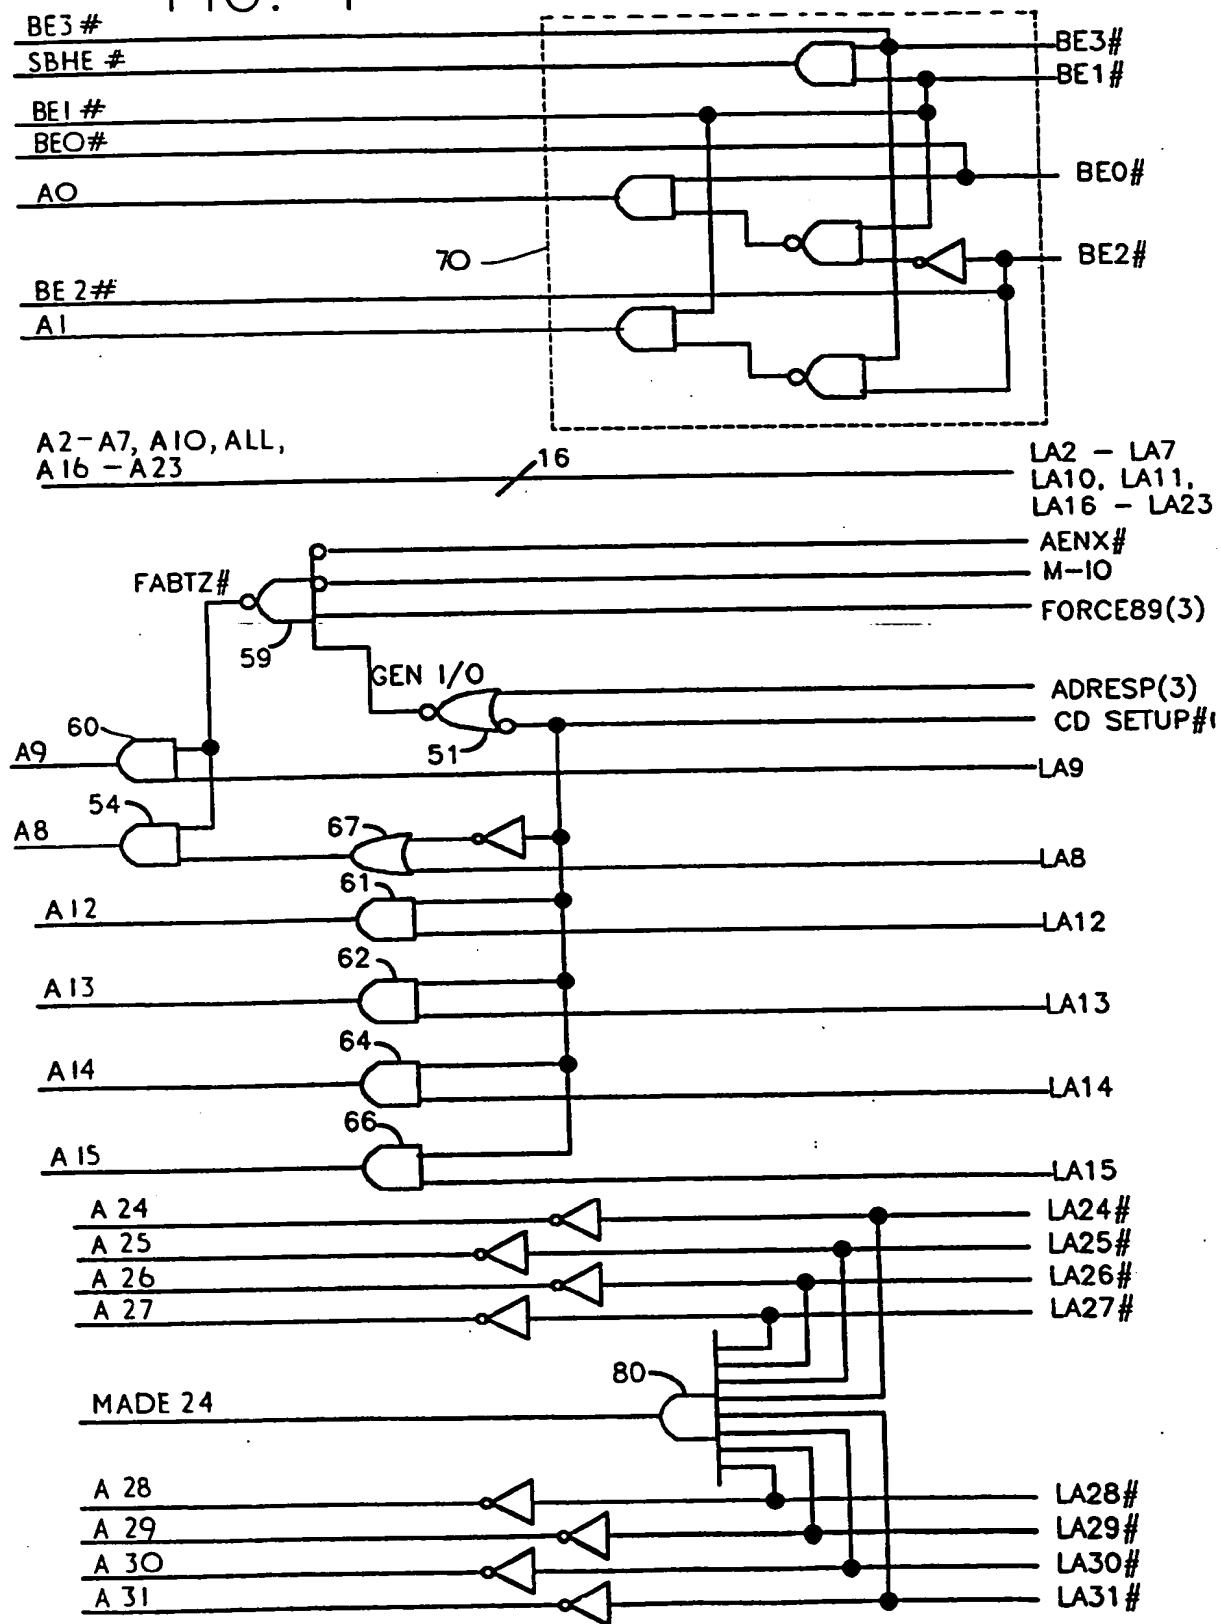


FIG. 3

FIG. 4



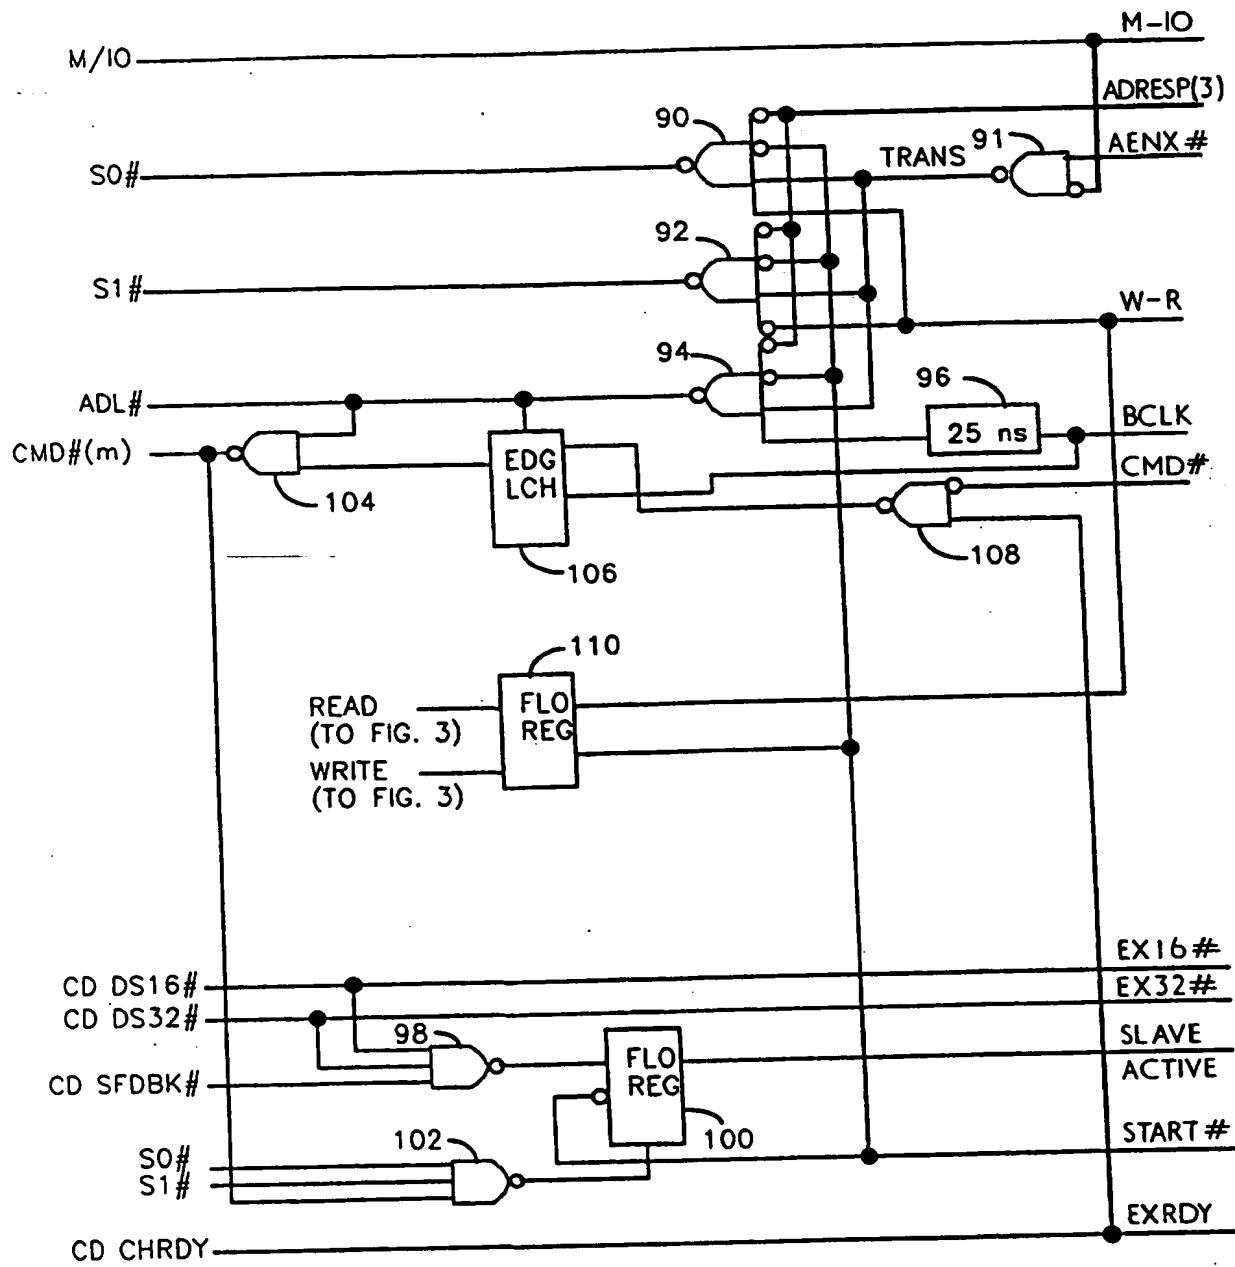


FIG. 5

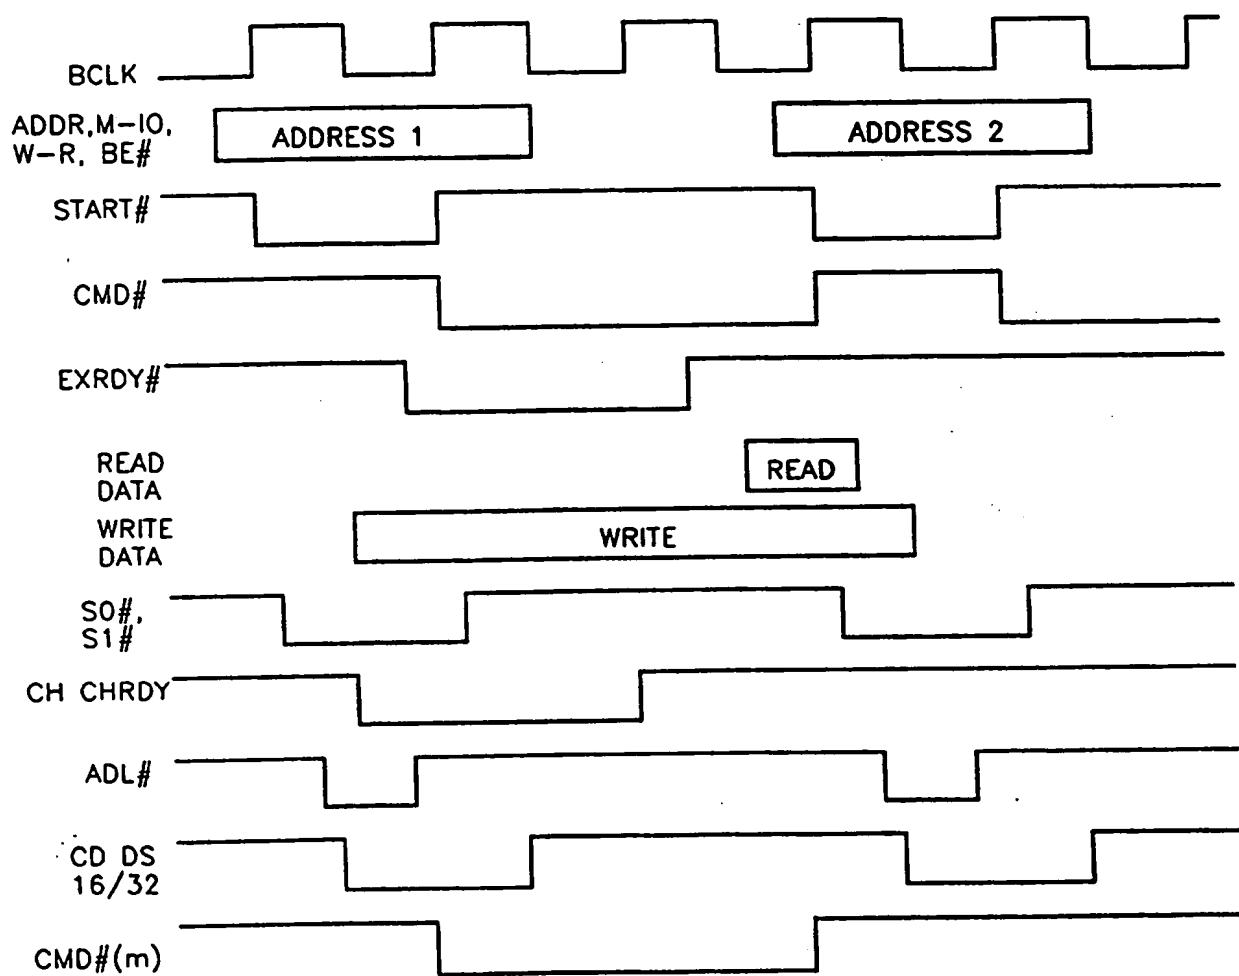


FIG. 6



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



0 488 501 A3

11 Publication number:

12

## EUROPEAN PATENT APPLICATION

21 Application number: 91308228.5

51 Int. Cl. 5. G06F 9/46, G06F 3/06 -

22 Date of filing: 09.09.91

23 Priority: 26.10.90 US 605054

71 Applicant: International Business Machines Corporation  
Old Orchard Road  
Armonk, N.Y. 10504(US)

43 Date of publication of application:  
03.06.92 Bulletin 92/23

72 Inventor: Miro, Felix  
4928 S.W. 90th Avenue  
Cooper City, Florida 33328(US)

84 Designated Contracting States:  
DE FR GB IT

74 Representative: Blakemore, Frederick Norman  
IBM United Kingdom Limited Intellectual  
Property Department Hursley Park  
Winchester Hampshire SO21 2JN (GB)

86 Date of deferred publication of the search report:  
21.04.93 Bulletin 93/16

### 54 Scheduling input/output operations in multitasking systems.

57 In a multitasking data processing system, I/O requests to a disk drive are staged in holding queues from which they are transferred to a service queue. Requests in the latter queue are directly serviced on a FIFO basis by a device driver module running on the system. The system maintains a set of holding queues and an associated service queue separately for each physical drive in the system (hard file, floppy drive, etc.). Holding queues in each set are prioritised in accordance with base priorities of tasks, and I/O requests to disk drives are entered into associated holding queues having priorities corresponding to those of task threads for which such requests are originated. Prioritisation of the holding queues, and a starvation advancement process performed to advance "oldest" enqueued requests to higher priority holding queues, causes the requests to be presented to the disk drive in a sequence based in part on respective task priorities and in part on "fairness" servicing of "service starved" requests. A selection operation in respect to certain transfers from the holding queues to the service

queues orders selection of transferred requests on the basis of relative proximity of disk addresses in the queued requests to current positions of the read/write head assembly in the associated disk drive. Forgoing operations serve to improve processing throughput of all tasks. A service kernel of the operating system, which maintains the holding queues, places each newly issued disk I/O request in a selected holding queue associated with a designated disk drive. The selected queue is one having a priority associated with the base priority of the task thread for which the request was issued. A request issued for a foreground task may be placed in a higher priority holding queue than requests of like priority issued for other tasks. When entering a new request into a holding queue, the kernel performs a "starvation" check relative to selected holding queues. If a starvation condition is detected, the kernel performs the above-mentioned starvation advancement operation to transfer requests in the affected queue to a higher priority queue.

EP 0 488 501 A3

EP 0 488 501 A3

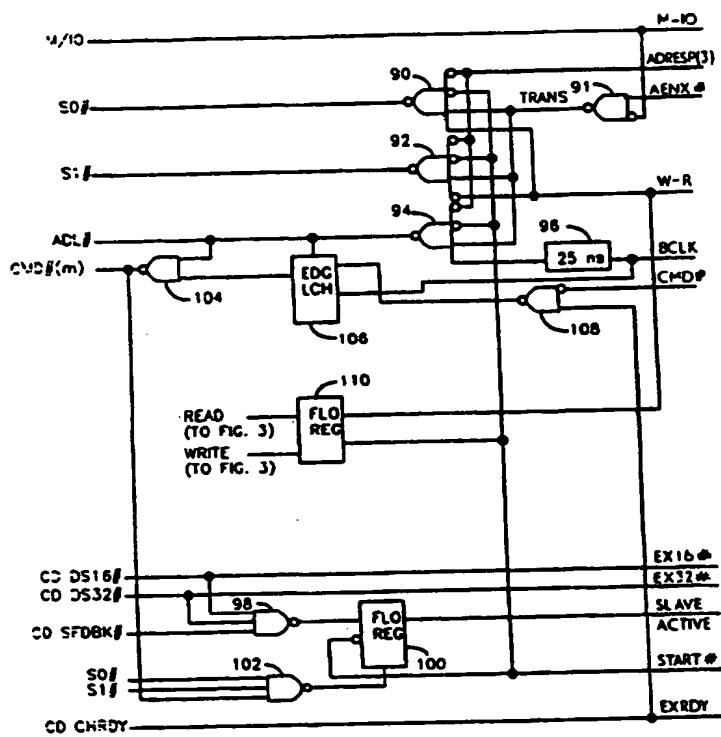


FIG. 5



European Patent  
Office

EUROPEAN SEARCH REPORT

Application Number

EP 91 30 8228  
Page 1

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int. CL.S)		
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. CL.S)		
Y	GB-A-2 181 871 (CODEX CORPORATION) * abstract * * figure 2 * * page 1, left column, line 1 - page 2, left column, line 8 * * claims 1,3,5-7,10,12 *	1-4,8-10	G06F9/46 G06F3/06		
A	---	5,11			
Y	PATENT ABSTRACTS OF JAPAN vol. 11, no. 207 (P-592)4 July 1987 & JP-A-62 027 818 (NEC CORPORATION) * abstract *	1-4,8-10			
A	---	7,12			
Y	IBM TECHNICAL DISCLOSURE BULLETIN vol. 25, no. 4, September 1982, NEW YORK US pages 2079 - 2080 A.H.NOLTA 'Batching DASD Accesses in a Multi-User Environment' * page 2079, line 5 - line 6 * * page 2080, line 4 - line 6 * * page 2080, line 14 - line 23 *	8	TECHNICAL FIELDS SEARCHED (Int. CL.S)		
A	---	7,12	G06F		
Y	IBM TECHNICAL DISCLOSURE BULLETIN vol. 27, no. 10B, March 1985, NEW YORK US pages 5967 - 5968 'Virtual Machine Group Control System Dispatching Program' * page 5967 *	8			
A	---	5-6,9,12			
The present search report has been drawn up for all claims					
Place of search	Date of completion of the search	Examiner			
THE HAGUE	09 FEBRUARY 1993	SCHARFENBERGER B.			
CATEGORY OF CITED DOCUMENTS					
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document					
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons					
A : member of the same patent family, corresponding document					



European Patent  
Office

EUROPEAN SEARCH REPORT

Application Number

EP 91 30 8228  
Page 2

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int. CL.5)						
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim							
Y	<p>IBM TECHNICAL DISCLOSURE BULLETIN vol. 31, no. 12, May 1989, NEW YORK US pages 271 - 272 'Priority Scheduling Algorithm' * page 271, line 1 - line 5 * * page 272, line 19 - line 30 * * figure 1 *</p> <p>-----</p>	2-4,8-10							
			TECHNICAL FIELDS SEARCHED (Int. CL.5)						
<p>The present search report has been drawn up for all claims</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">Place of search</td> <td style="width: 33%;">Date of completion of the search</td> <td style="width: 34%;">Examiner</td> </tr> <tr> <td>THE HAGUE</td> <td>09 FEBRUARY 1993</td> <td>SCHARFENBERGER B.</td> </tr> </table> <p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document</p>				Place of search	Date of completion of the search	Examiner	THE HAGUE	09 FEBRUARY 1993	SCHARFENBERGER B.
Place of search	Date of completion of the search	Examiner							
THE HAGUE	09 FEBRUARY 1993	SCHARFENBERGER B.							